

METHOD TO PROVIDE SYNC NOTIFICATIONS TO CLIENT DEVICES

FIELD OF THE INVENTION

[0001] This invention relates generally to computer based notification systems and, more particularly, relates to a method and system to notify a user of events of interest to the user.

BACKGROUND OF THE INVENTION

[0002] With the development and deployment of mobile devices such as laptop computers, personal digital assistant devices, cell phones, etc. and infrastructures, consumers and businesses are increasingly being able to realize the benefits of true mobile computing, collaboration, and information exchange. No longer are business travelers required to carry an assortment of cables and search endlessly for an available data port simply to connect to a network to retrieve email messages, download files, or exchange information. No longer are companies and home consumers restrained in where they may access their networks by the location of the Ethernet jacks on the wall. Meeting participants and groups of friends may now form their own ad hoc networks without connecting cables between themselves or logging in to some preexisting network. They can log onto the network using a wireless protocol while running on battery power, thereby allowing even greater mobility

[0003] Hand-in-hand with mobile devices, electronic messaging such as electronic mail ("e-mail") and instant messaging carried over the Internet, is rapidly becoming not

only quite pervasive in society but also, given its informality, ease of use and low cost, a preferred method of communication for many individuals and organizations. However, electronic messaging requires the users of the mobile devices to be within the range of the wireless networks they are on or subscribe to in order to receive messages. When users come back in the network after being out of range or when first logging in, the users often want to know if there are any messages waiting. Additionally, many users want to be notified when an event of interest occurs such as a particular message arriving, news of an event (e.g., calendar update), etc. so that the users can connect with the server of the network and retrieve the message/event/etc.

[0004] The most straightforward way to ensure that a user is kept in sync with a server is to send a sync notification (e.g., a ping) to the client whenever an event of interest occurs on the server. The client then connects back to the server to retrieve the data. This methodology requires a large number of notifications, which in turn consumes additional server CPU cycles, network bandwidth and battery life on the client. In today's network environments where users of a network are charged based upon time of day, length of time, and/or bandwidth usage, the notifications can increase the cost for the user. For example, where notifications are delivered over SMS (Short Messaging Service), this approach can significantly increase a user's monthly bill and waste carrier resources.

[0005] Industry has taken some steps to mitigate the cost of notifications. Some mitigations include waiting for a certain amount of time to elapse or a certain number of

items to accumulate before sending the notification. However, all such approaches introduce latency into the system that does not have to be there. These approaches make a user wait five to ten minutes or longer for their email when the user can get it right away.

BRIEF SUMMARY OF THE INVENTION

[0006] The invention maintains a notion of the state of the mobile device, indicating whether it is up to date or not from the perspective of the server. When an event occurs that is of interest to the mobile device, the server refers to the stored state to decide whether or not the device needs to be prompted to sync.

[0007] The server receives notification that an event of interest has been received. A parameter (e.g., a syncGUID) that indicates whether an event of interest renders the device no longer up to date from the perspective of the server or whether an event of interest occurs and the server has not been contacted by the device for a certain period of time is compared to a last known value of the parameter for the device. The parameter is stored in a configuration file for the user/device.

[0008] If the parameter matches the last known value of the parameter that is stored in memory and the current time is greater than a timeout value, a sync notification is sent and a new timeout value is set. The sync notification may be sent using the SMTP (simple mail transfer protocol) protocol. If the parameter does not match the last known

value, the last known value is set to the parameter value and a sync notification is sent.

The timeout value is also set to a new value.

[0009] The new value for the timeout is set to be equal to the current time plus a predetermined value. The predetermined value is a parameter set by the server and in one embodiment, is typically set based on the confidence level of the network and is how long to wait to retry sending the notification to the device. The value can be set to any value. A value of one to two hours is sufficient for many networks while a value of fifteen minutes is needed for other networks.

[0010] Additional features and advantages of the invention will be made apparent from the following detailed description of illustrative embodiments which proceeds with reference to the accompanying figures.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] While the appended claims set forth the features of the present invention with particularity, the invention, together with its objects and advantages, may be best understood from the following detailed description taken in conjunction with the accompanying drawings of which:

[0012] FIG. 1 is a block diagram generally illustrating an exemplary computer system on which the present invention resides;

[0013] FIG. 2 is a block diagram generally illustrating an exemplary environment in which the present invention may operate;

[0014] FIG. 3 is a flow chart illustrating a process for providing a sync notification to a client; and

[0015] Figure 4 is a block diagram illustrating the communication flow of the process of FIG 3 in the environment of FIG. 2.

DETAILED DESCRIPTION OF THE INVENTION

[0016] The invention provides a process for notifying a user of events of interest via sync notifications if the user's device has not been synchronized with a server. The invention maintains a notion of the state of the mobile device, indicating whether it is up to date or not from the perspective of the server. When an event occurs that is of interest to the mobile device, the server refers to the stored state to decide whether or not the device needs to be prompted to sync. The mobile device is prompted to sync if the server decides that the device needs to be prompted to sync. Using this approach, common scenarios such as a user turning off her mobile device for a few hours will require far fewer notifications to maintain the client. For example, a user turns off her mobile device for a few hours while she's out to lunch. During that period of time, she gets twenty-five e-mails. As soon as she turns her device back on, she'll receive one sync notification and her device will be up to date within minutes. Conventional notification systems would

send twenty-five notifications or wait a period of time before sending out at least one notification.

[0017] The invention may be implemented in a system employing various types of machines, including cell phones, hand-held devices, wireless surveillance devices, microprocessor-based programmable consumer electronics, and the like, using instructions, such as program modules, that are executed by a processor. Generally, program modules include routines, objects, components, data structures and the like that perform particular tasks or implement particular abstract data types. The term “program” includes one or more program modules.

[0018] Turning to the drawings, wherein like reference numerals refer to like elements, the invention is illustrated as being implemented in a suitable computing environment. Although not required, the invention will be described in the general context of computer-executable instructions, such as program modules, being executed by a personal computer. Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. Moreover, those skilled in the art will appreciate that the invention may be practiced with other computer system configurations, including hand-held devices, multi-processor systems, microprocessor based or programmable consumer electronics, network PCs, minicomputers, mainframe computers, and the like. The invention may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications

network. In a distributed computing environment, program modules may be located in both local and remote memory storage devices.

[0019] FIG. 1 shows an exemplary computing device 100 for implementing an embodiment of the invention. In its most basic configuration, the computing device 100 includes at least a processing unit 102 and a memory 104. Depending on the exact configuration and type of computing device, the memory 104 may be volatile (such as RAM), non-volatile (such as ROM, flash memory, etc.) or some combination of the two. This most basic configuration is illustrated in FIG.1 by a dashed line 106. Additionally, the device 100 may also have additional features/functionality. For example, the device 100 may also include additional storage (removable and/or non-removable) including, but not limited to, magnetic or optical disks or tapes. Such additional storage is illustrated in FIG. 1 by a removable storage 108 and a non-removable storage 110. Computer storage media includes volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information such as computer readable instructions, data structures, program modules or other data. The memory 104, the removable storage 108 and the non-removable storage 110 are all examples of computer storage media. Computer storage media includes, but is not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CDROM, digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can accessed by the device 100. Any such computer storage media may be part of the device 100.

[0020] The device 100 may also contain one or more communications connections 112 that allow the device to communicate with other devices. The communications connections 112 are an example of communication media. Communication media typically embodies computer readable instructions, data structures, program modules or other data in a modulated data signal such as a carrier wave or other transport mechanism and includes any information delivery media. The term “modulated data signal” means a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not limitation, communication media includes wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, RF, infrared and other wireless media. As discussed above, the term computer readable media as used herein includes both storage media and communication media.

[0021] The device 100 may also have one or more input devices 114 such as keyboard, mouse, pen, voice input device, touch-input device, etc. One or more output devices 116 such as a display, speakers, printer, etc. may also be included. All these devices are well known in the art and need not be discussed at greater length here.

[0022] The device 100 is configured as a mobile device. While the device 100 is configured as a mobile device, those skilled in the art will recognize that the device does not have to be a mobile device; it can be a desktop device. The device 100 is provided with a portable power source 120, such as a battery pack, a fuel cell, or the like. The

power source 120 provides power for computations and wireless data transmissions by the device 100. It should be noted that the power source 120 might further include an external power source that overrides or recharges the built-in batteries/fuel cell, such as an AC adapter or a powered docking cradle for connection to a wired network. The mobile device 100 may further include a network interface card (NIC) for wirelessly communicating with different types of wireless networks. The NIC includes a transmitter, which is coupled to an antenna for transmitting data wirelessly over a suitable frequency channel. A receiver is also coupled to the antenna for receiving communication packets wirelessly transmitted from the networks with which the device is communicating. The network interface module card and the antenna are part of the communication connections 112. It will be appreciated that the interface type and physical configuration of the network interface module is not critical to the invention. For instance, the interface type could be PCI or another type and the network interface module does not have to reside on a separate card. It may be included on the motherboard of the computer or even possibly built into the processor in the future.

[0023] Through the wireless network interface module, the mobile device 100 may communicate with different types of wireless networks. For instance, the mobile device 100 may be connected wirelessly to an infrastructure network through an access point thereof. The mobile device 100 may also be part of a peer-to-peer network, also referred to as an *ad hoc* network, which includes other wireless devices.

[0024] In the description that follows, the invention will be described with reference to acts and symbolic representations of operations that are performed by one or more computers, unless indicated otherwise. As such, it will be understood that such acts and operations, which are at times referred to as being computer-executed, include the manipulation by the processing unit of the computer of electrical signals representing data in a structured form. This manipulation transforms the data or maintains it at locations in the memory system of the computer, which reconfigures or otherwise alters the operation of the computer in a manner well understood by those skilled in the art. The data structures where data is maintained are physical locations of the memory that have particular properties defined by the format of the data. However, while the invention is being described in the foregoing context, it is not meant to be limiting as those of skill in the art will appreciate that various of the acts and operation described hereinafter may also be implemented in hardware.

[0025] An exemplary environment 200 in which the invention operates is shown in FIG 2. An Exchange e-mail server 202 shall be used to describe the invention. Exchange is an e-mail server produced by Microsoft Corporation. While Exchange will be used, the invention may be used with other messaging servers, including instant message servers. The Exchange server 202 communicates with devices 204, 206 on a wired network 208 and devices 210-216 in range of wireless network 218. As shown in figure 2, mobile device 216 is not in range of network 218.

[0026] Turning now to FIG. 3, the server 202 detects when an event of interest for a user has occurred. The event of interest can be an e-mail arriving, a new or edited calendar item, a response to an inquiry, a traffic update, a weather update, instant messaging, etc. The device/user configuration file is retrieved for the user (step 300). The device/user configuration file has the phone number (if applicable), user name, and other particulars about the user. In the device/configuration file is a parameter that provides a representation of the state of the device and it provides an indication of whether an event of interest renders the device no longer up to date from the perspective of the server or whether an event of interest occurs and the server has not been contacted by the device for a certain period of time.

[0027] The device is up-to-date when, to the server's knowledge, the device is completely in sync with the server (with the possible exception of the event which just triggered a sync notification as described below). The device is pending synchronization when a sync notification has been sent to the device telling it to sync with the server, but the device has not yet performed synchronization. The parameter that can be used is called a syncGUID, which is used primarily to prevent spoofing/replay attacks. The syncGUID happens to have the same semantics. The syncGUID for a device is updated after each successful device synchronization with the server. The device/user configuration file contains the syncGUID or a link to the syncGUID. As used herein, the syncGUID shall be used to designate a parameter that provides a representation of the state of the device and it provides an indication of whether an event of interest renders the device no longer up to date from the perspective of the server or whether an event of

interest occurs and the server has not been contacted by the device for a certain period of time.

[0028] The server maintains an in memory table of "last-known" syncGUIDs for each device, which for purposes of explanation shall be called a trackingGUID. If the syncGUID matches the trackingGUID for a given device, it can be concluded that the device has not performed a synchronization since the last event/notification was processed. This means that the device is in the pending synchronization state. If the GUIDs don't match, we conclude that the device has performed a sync since the last event/notification was processed, indicating that it is in the up-to-date state prior to the event that just triggered the notification.

[0029] The server checks to see if the user has retrieved the message. If the user has retrieved the message, no sync notification is sent. If the user has not retrieved the message, the following steps are performed. For each device listed in the user's configuration file, the trackingGUID is compared to the syncGUID. If the trackingGUID is not equal to the syncGUID, the trackingGUID is set equal to the syncGUID (step 302). The timeout value is set to the current time plus X minutes (step 304). X minutes is a parameter set by the server and in one embodiment, is typically set based on the confidence level of the network 208, 218. For example, if a network is not reliable and a message gets dropped, the X parameter is how long to wait to retry sending the notification to the device. The value of the X parameter can be set to any value. A value of one to two hours is sufficient for many networks while a value of fifteen minutes is

needed for other networks. Once the timeout value is set, the sync notification is sent to the device (step 306).

[0030] If the trackingGUID is equal to the syncGUID, the current time is compared to the timeout time. If the time is less than the timeout value, no sync notification is sent. If the time is greater than the timeout value, a sync notification is sent (step 308) and a new timeout value is set (step 310). The timeout value is set by setting the timeout value to the current time plus X minutes as described above. Note that while FIG. 3 shows waiting for another event to occur if the time is less than the timeout value, an alternate approach is to determine if the user has retrieved the message when the time is greater than or equal to the timeout value and then perform steps 308 and 310 if the user has not retrieved the message during the interim time.

[0031] Turning now to FIG. 4, an implementation of the process is shown using the Exchange server 202 and device 216 of FIG 2. An event of interest such as an email arrives at the server 202 and put in the store 402. The store represents any information store that supports eventing. The arrival of the event creates a trigger 404 that is sent to module 406. The module 406 performs the process of steps 300 to 310 after receiving the trigger 404. If applicable, a sync notification 408 is sent to device 216 via SMS. Following the process of steps 300 to 310, notifications 408 are sent to device 216 until the device 216 syncs. The Exchange server 202 has an air sync protocol 410 that devices (e.g., pocket PCs and phones) talk to in order to retrieve data from the store 402 that the devices and module 406 read from and write to in order to send and retrieve data. Once

the device 216 syncs with the server as represented by line 412, the error sync protocol updates the syncGUID to the device/user configuration file as represented by line 414.

[0032] The type of sync notification may be a text message or it can be based on the device type. If a device supports image rendering, the sync notification can take the form of a graphical image or a text image while a device that supports only text rendering would only receive text based sync notifications. For example, if the event of interest is a daily summary of a user's calendar, a graphical image of the summary is sent if the device supports graphical rendering. A text summary is sent if the device only supports text rendering.

[0033] In one embodiment, user settable filters are used to determine when the sync notification is to be sent. These filters can be set up to send certain types of notifications right away and other types after the next timeout event, send notifications only during business hours or only during off-hours, and the like. Module 406 sends the sync notifications in accordance with the filter settings. The filter settings can be stored in the device/user configuration file or in a separate file.

[0034] It can be seen from the foregoing that a method to provide sync notifications to a device has been described. The invention maintains a notion of the state of the device. When an event of interest occurs for the mobile device, the server decides whether or not the device needs to be prompted to sync. The device is prompted to sync if the server decides that the device needs to be prompted to sync. This method makes

more efficient use of network resources than conventional systems and provides lower end user cost without introducing additional latency into the system.

[0035] In view of the many possible embodiments to which the principles of this invention may be applied, it should be recognized that the embodiment described herein with respect to the drawing figures is meant to be illustrative only and should not be taken as limiting the scope of invention. For example, those of skill in the art will recognize that the elements of the illustrated embodiment shown in software may be implemented in hardware and vice versa or that the illustrated embodiment can be modified in arrangement and detail without departing from the spirit of the invention. Therefore, the invention as described herein contemplates all such embodiments as may come within the scope of the following claims and equivalents thereof.